

IN THE UNITED STATES DISTRICT COURT  
FOR THE NORTHERN DISTRICT OF CALIFORNIA

SUN MICROSYSTEMS, INC.,

No. C-07-05488 EDL

Plaintiff,

v.

**ORDER GRANTING SUMMARY  
JUDGMENT OF NON-INFRINGEMENT  
OF U.S. PATENT NO. 5,124,987**

NETWORK APPLIANCE, INC.,

Defendant.

---

**I. INTRODUCTION**

On October 29, 2007, Sun Microsystems, Inc. (“Sun”) filed its Complaint, alleging that Network Appliance, Inc. (“NetApp”) infringed and is infringing, directly and indirectly under 35 U.S.C. § 271, certain of its patents, by making, using, selling, or offering for sale certain data processing systems and related software. Sun seeks a declaratory judgment that certain patents owned by Sun are each not infringed, are invalid or unenforceable, as well as a permanent injunction and damages. On December 21, 2007, NetApp filed an Answer and Counterclaim, denying the material allegations of Sun’s Complaint and asserting a number of affirmative defenses and counterclaims. NetApp denies infringing any of the Sun Patents, including the patent at issue in this motion (U.S. Patent Number 5,124,987 (the “‘987 Patent”)) and alleges that Sun infringes a number of its patents instead. On December 22, 2008, this Court issued an Order Construing Claims (the “12/22/08 Order”) in which it construed certain disputed terms and phrases contained in various claims in the patents at issue between the parties, including one terms contained in the ‘987 patent. The parties subsequently conducted discovery, and each party has filed two motions in the above-captioned 07-5488 case.

On December 2, 2009, NetApp filed a Motion For Summary Judgment Of Non-Infringement Of U.S. Patent No. 5,124,987 (the “‘987 Motion”) on the basis that its allegedly infringing product, WAFL, does not practice the “first available memory space” claim limitation of the ‘987 patent. NetApp also argues that each asserted claim requires that modified data be written “in response to” the subsequent receipt of modifications to one of said data records “exclusive of” the rest of that received stream of data records, and that WAFL also does not practice this aspect of the invention. The ‘987 Motion was fully briefed, and a hearing was held on January 20, 2010. Having considered the record in this case and the parties’ statements at oral argument, and for the reasons set forth below, the Court hereby GRANTS NetApp’s Motion For Summary Judgment Of Non-Infringement of the ‘987 patent.

## II. LEGAL STANDARD

### A. Summary Judgment

Summary judgment shall be granted if “the pleadings, discovery and disclosure materials on file, and any affidavits show that there is no genuine issue as to any material fact and that the movant is entitled to judgment as a matter of law.” Fed. R. Civ. Pro. 56(c). Material facts are those which may affect the outcome of the case. See Anderson v. Liberty Lobby, Inc., 477 U.S. 242, 248 (1986). A dispute as to a material fact is genuine if there is sufficient evidence for a reasonable jury to return a verdict for the nonmoving party. Id. The court must view the facts in the light most favorable to the non-moving party and give it the benefit of all reasonable inferences to be drawn from those facts. Matsushita Elec. Indus. Co. v. Zenith Radio Corp., 475 U.S. 574, 587 (1986). The court must not weigh the evidence or determine the truth of the matter, but only determine whether there is a genuine issue for trial. Balint v. Carson City, 180 F.3d 1047, 1054 (9th Cir. 1999).

A party seeking summary judgment bears the initial burden of informing the court of the basis for its motion, and of identifying those portions of the pleadings and discovery responses that demonstrate the absence of a genuine issue of material fact. Celotex Corp. v. Catrett, 477 U.S. 317, 323 (1986). Where the moving party will have the burden of proof at trial, it must affirmatively demonstrate that no reasonable trier of fact could find other than for the moving party. On an issue where the nonmoving party will bear the burden of proof at trial, the moving party can prevail

merely by pointing out to the district court that there is an absence of evidence to support the nonmoving party's case. Id. If the moving party meets its initial burden, the opposing party "may not rely merely on allegations or denials in its own pleading;" rather, it must set forth "specific facts showing a genuine issue for trial." See Fed. R. Civ. P. 56(e)(2); Anderson, 477 U.S. at 250. If the nonmoving party fails to show that there is a genuine issue for trial, "the moving party is entitled to judgment as a matter of law." Celotex, 477 U.S. at 323.

### **B. Patent Infringement**

"To prove infringement, the patentee must show that the accused device meets each claim limitation either literally or under the doctrine of equivalents." Catalina Mktg. Int'l v. Coolsavings.com, Inc., 289 F.3d 801, 812 (Fed. Cir. 2002). A determination of infringement, whether literal or under the doctrine of equivalents, is a question of fact. Id. "Literal infringement requires the patentee to prove that the accused device contains each limitation of the asserted claim." Id. "Summary judgment of no literal infringement is proper when, construing the facts in a manner most favorable to the nonmovant, no reasonable jury could find that the accused system meets every limitation recited in the properly construed claims." Id. Where the parties do not dispute any relevant facts regarding the accused product, but disagree over possible claim interpretations, the question of literal infringement collapses into claim construction and is amenable to summary judgment. General Mills, Inc. v. Hunt-Wesson, Inc., 103 F.3d 978, 983 (Fed. Cir. 1997); cf. Int'l Rectifier Corp. v. IXYS Corp., 361 F.3d 1363, 1375 (Fed. Cir. 2004) (distinguishing General Mills on the basis that only the structure of the accused devices had been stipulated to, not the disputed factual determination of whether the device met the claims as construed, but not addressing the scenario in which no reasonable juror could find that a certain claim limitation was met).

In MyMail Ltd. v. America Online, Inc., 476 F.3d 1372, 1378 (Fed. Cir. 2007), the Federal Circuit reviewed a District Court order granting summary judgment of non-infringement. Because there were no material factual disputes as to the operation of the accused systems, and the parties' disagreements concerned whether the defendants' systems performed "authentication" set forth in the patent and construed by the district court, the Federal Circuit held that the issue reduced to a question of claim interpretation and affirmed summary judgment. See id. (noting that the accused

product did not satisfy the authentication requirement as it did not validate the user's ID and password, as required by the patent's authentication process). These cases teach that the Court cannot leave it to the jury to decide the proper scope of the patent claim terms. 02 Micro Int'l Ltd. v. Beyond Innovation Tech. Co. Ltd., 521 F.3d 1351, 1360 (Fed. Cir. 2008) ("When the parties raise an actual dispute regarding the proper scope of the[] claims, the court, not the jury, must resolve the dispute.").

"Infringement under the doctrine of equivalents requires the patentee to prove that the accused device contains an equivalent for each limitation not literally satisfied." Id. The Court may not apply the doctrine of equivalents so as to vitiate a claim limitation. Warner-Jenkinson, 520 U.S. at 29, 39 n.8. The Federal Circuit articulates the test for equivalence in two different ways. See Voda v. Cordis Corp., 536 F.3d 1311, 1326 (Fed. Cir. 2008). Under the insubstantial differences test, "[a]n element in the accused device is equivalent to a claim limitation if the only differences between the two are insubstantial." Honeywell Int'l Inc. v. Hamilton Sundstrand Corp., 370 F.3d 1131, 1139 (Fed.Cir. 2004); Warner-Jenkinson Co. v. Hilton Davis Chem. Co., 520 U.S. 17, 40 (1997)). Alternatively, under the function-way-result test, an element in the accused device is equivalent to a claim limitation if it "performs substantially the same function in substantially the same way to obtain substantially the same result." Schoell v. Regal Marine Indus., Inc., 247 F.3d 1202, 1209-10 (Fed. Cir. 2001). "Where the evidence is such that no reasonable jury could determine two elements to be equivalent," summary judgment of non-infringement under the doctrine of equivalents is proper. Warner-Jenkinson, 520 U.S. at 39 n. 8. Summary judgment has been rejected because of conflicting expert testimony on the application of the function-way-result test. Crown Packaging Tech., Inc. v. Rexam Bev. Can Co., 559 F.3d 1308, 1315 (Fed. Cir. 2009) (holding that conflicting expert evidence regarding function establishes material issue of fact).

### **III. NETAPP'S MOTION FOR SUMMARY JUDGMENT OF NON-INFRINGEMENT OF U.S. PATENT NO. 5,124,987**

#### **A. Patent and Claim Construction Background**

The '987 patent, "Logical Track Write Scheduling System for a Parallel Disk Drive Array Data Storage Subsystem," is directed to a data storage subsystem. The data storage subsystem stores

1 data on a number of small hard disk drives, but emulates the format and operation of a large form  
 2 factor disk drive: “This invention relates to an inexpensive, high performance, high reliability  
 3 parallel disk drive array data storage subsystem that includes an efficient data storage management  
 4 system to dynamically map virtual data storage devices to logical data storage devices and schedule  
 5 the writing of data on these devices.” See Ho Decl. Ex. A (‘987 Patent) at 1:7-12. The independent  
 6 claims of the ‘987 Patent asserted by Sun are 9, 57, and 73. Each asserted claim requires selecting  
 7 and writing data to a “first available memory space” (claims 9 and 57) or a “first available logical  
 8 track” (claim 73). In addition, each asserted claim requires that modified data is written “in  
 9 response to” the subsequent receipt of modifications to one of said data records “exclusive of” the  
 10 rest of said received stream of data records.

11 For example, claim 9 provides:

12 In a disk memory system, having a plurality of disk drives, a number of said  
 13 plurality of disk drives being configured into at least two redundancy groups, each  
 14 redundancy group consisting of at least two disk drives, a method of storing data  
 15 records for at least one associated data processor comprising the steps of:

16 selecting, in response to the receipt of a stream of data records from said  
 17 associated data processor, **first available memory space** in one of said  
 18 redundancy groups to store said received stream of data records thereon;

19 writing said received stream of data records and redundancy data associated  
 20 with said received stream of data records in said selected **first available**  
 21 **memory space** in said selected redundancy group;

22 writing, **in response to** the subsequent receipt of modifications to one of  
 23 said data records stored in said **first available memory space** in said  
 24 selected redundancy group from said associated data processor, said  
 25 modified data record, **exclusive of** the rest of said received stream of data  
 26 records and said redundancy data associated with said received stream of  
 27 data records written in said **first available memory space**, in second  
 28 available memory space in one of said redundancy groups by including said  
 modified data record with a stream of data records subsequently received for  
 said step of writing; and

converting said **first available memory space** used to store said originally  
 received data record to available memory space.

1 In its Claim Construction Order, the Court adopted NetApp’s construction of “first available  
 2 memory space” as “one or more logical tracks, each of which must be empty, i.e., recognized by the  
 3 system as available free space.” 12/22/08 Claim Construction Order at 10. The Court noted Sun’s  
 4 concession that, “whenever claim 9 or 57 is practiced, those claims require writing to an empty

logical track,” but declined to rule at that time on the issue of whether an accused system “that fully practices the claimed method most of the time nevertheless does not infringe the patent if it ever practiced a different method.” Id. However, the Court observed that “there is strong support for [this position] in the specification and prosecution history.” Id. at 9. The Court stated that, “while the applicants’ statements in the prosecution history arguably do not by themselves amount to a clear and unambiguous statement that a system must always practice the claimed method in order to practice the patent, they lend further support to the strong language in the specification describing the feature of the invention as a whole, and thus may well limit the scope of the invention to require always writing only to an empty track.” Id. at 10.

The Court has not construed the term “in response to” as it is used in the ‘987 Patent. However, the Court has construed this term in connection with Sun’s Patent Number 5,549,857 (the “‘857 Patent”), where it was construed to mean “after and in reaction to.” See Ho Decl., Ex. F (9/10/08 Claim Construction Order in related case 07-6053). In an Order granting summary judgment of non-infringement with respect to the ‘857 Patent, the Court noted that its construction of the term “intended to clarify the requirement of causation in the process: the writing must be a real, cause-in-fact trigger for the transmission.” 11/17/09 Order Granting Summary Judgment of Non-Infringement of ‘857 Patent at 10. The parties dispute the applicability of this prior construction to the instant motion. The Court has not construed the term “exclusive of,” but to date there has been little or no dispute about its meaning.

#### **B. Sun’s Infringement Assertions**

Sun’s infringement assertions concern NetApp’s storage system products running DataONTAP, WAFL and associated RAID 4/DP software and hardware, but the summary judgment arguments are directed primarily to WAFL because it is included in all of the accused technologies. Sun argues that WAFL literally infringes the ‘987 patent because it is optimized to write full stripes to empty logical tracks most or all of the time, and this is sufficient to establish infringement. In its Opposition, it points to evidence that WAFL writes to empty logical tracks at least 97-99% of the time. In contrast, NetApp contends that WAFL’s write allocation process is “agnostic” with respect to whether the RAID stripes it allocates for writing data are empty, and the system instead targets

empty blocks within a particular allocation area of a RAID group that may or may not include empty stripes. Motion at 8:9-15 (citing Ho Decl., Ex. I (Mattson Expert Report) at ¶ 60).<sup>1</sup>

### C. NetApp's Motion

NetApp's Motion argues that there is no triable issue of fact as to the "first available memory space" limitation in the asserted claims. Specifically, NetApp focuses on the fact that the '987 Patent was intended to solve a "performance bottleneck" problem associated with previous RAID systems, which needed to update parity information in a redundancy group every time one of the data files in the group was modified. See Ho Decl., Ex. C ('987 Prosecution History, 9/23/91 Amendment) at NAB0013338 (stating that the need for repeated parity updates has a negative impact on performance, because extra processing time is needed to ensure that parity data for a redundancy group is up to date whenever data is modified); see also Ex. A ('987 Patent) 2:67-3:23 ("A significant difficulty with the Fourth and Fifth Level RAID systems is that the parity for the parity group must be updated every time data is written into the parity group"). NetApp contends that the patent specifications and prosecution history of the '987 Patent make clear that the patent is limited to a method of storing data records where parity data is *never* updated because new or modified data is *always* written to empty logical tracks. Motion at 2:20-5:5; see also '987 Patent at 3:26-37 ("This system avoids the parity update problem of the prior art by never updating the parity in a data redundancy group. Instead all new or modified data is written on empty logical tracks and the old data is tagged as obsolete."). Because WAFL does not *always* select or write to empty RAID stripes (though NetApp admits that it may do so frequently as a side-effect of selecting under-allocated areas), and therefore sometimes updates parity data in place, NetApp argues that there is no infringement. As discussed below, this question essentially boils down to the one presented, but not resolved, during claim construction – whether or not the '987 patent requires that an accused

---

<sup>1</sup>Neither party really addresses the difference, if any, between a "logical track" and a "RAID stripe" or discusses whether this difference is important to the Court's analysis. NetApp states that a "stripe" is a set of locations on the disks of a redundancy group where data and redundancy data is stored and a "logical track" is defined by the patent as "the set of all physical tracks in a redundancy group which have the same physical track address." Ho Decl. Ex. A at 14:40-42. Sun appears to use the terms interchangeably, and notes that writing a full stripe of data to a logical track requires that the logical track be empty because a full stripe requires the availability of all of the blocks in the logical track. Opp. at n.1. The parties dispute whether these terms are synonymous, but agree that the Court need not reach the issue to grant summary judgment.



1 technology *always* practice the method of claims to infringe, or whether it is sufficient if the accused  
2 technology practices the method *most of the time*.

3 Additionally, NetApp argues that the '987 patent requires that one data record be written to a  
4 second available memory space "in response to" the subsequent receipt of modifications to one of  
5 said data records stored in a first available memory space. NetApp interprets this to mean that the  
6 writing step must be performed with each modification of a previously written data record and the  
7 receipt of a modified data record must be the cause-in-fact trigger for the writing step. NetApp  
8 argues that WAFL does not infringe this technology because it does not act upon each write request  
9 as it is received by the filer, but instead temporarily stores new or modified data in memory and  
10 flushes batches to disk when performing a consistency point. Finally, NetApp argues that there is no  
11 evidence that WAFL writes only *one* modified data record, "exclusive of" the rest of a stream of  
12 data records and other modified records from the same stream, and for this additional reason each  
13 element of the claims has not been not met.

#### 14 **D. The "First Available Memory Space" Limitation**

##### 15 **1. The '987 Patent Requires that a System Always Writes to Empty 16 Logical Tracks**

17 NetApp contends that the '987 Patent requires that an accused system *always* write to empty  
18 logical tracks as set forth in the specification and file history in order to solve the "parity update  
19 problem," i.e., the extra processing time needed to ensure that the parity data for a redundancy group  
20 is up to date whenever data is modified. See Ho Decl., Ex. A at 2:67-3:23; Ex. C at NAB0013338  
21 ('987 Patent File History, 9/17/91 Amendment) ("It is this exact problem that Applicants' invention  
22 solves by never updating parity in any redundancy group at any time."), NAB0013346 ("Applicants  
23 believe that the combined use of Applicants' apparatus, wherein data records are never updated in  
24 place..."); Ex. E (Smith Depo.) at 329-330.<sup>2</sup> NetApp points to Sun's response to a PTO rejection on  
25

---

26 <sup>2</sup> NetApp also relies on the testimony of inventor George Rudeseal, who explained that all of  
27 the data that is written is written to empty logical tracks. See Ho Decl. Ex. D at 358, 361, 404. Sun  
28 points out that this Court has previously followed federal circuit authority holding that "inventor  
testimony concerning the scope of the claims is irrelevant to [claim] construction." See Sun  
Microsystems, Inc. v. Network Appliance, Inc., 2009 WL 1513384 at \*14 (N.D. Cal. 2009) (citing  
Howmedica Osteonics Corp. v. Wright Med. Techn., Inc., 540 F.3d 1337, 1346-47 n. 5 (Fed.Cir.2008)).



the basis that it was “unclear why writing the data to a new location is an advantage,” see id. Ex. C at NAB0013334, in which Sun stated that prior systems were highly inefficient and that its invention solved the problem by “*never* updating parity in any redundancy group at any time.” Id. at NAB0013338 (emphasis added). In this response to the PTO, Sun further stated that “modified data, as well as all new data, *all flows to unused storage space* on the disk drives.” Id. (emphasis added).

NetApp also points to the “Solution” section of the specification, which explains that: “This system avoids the parity update problem of the prior art by *never* updating the parity in a data redundancy group. Instead, *all new or modified data is written on empty logical tracks* and the old data is tagged as obsolete.” Ex. A (‘987 Patent) at 3:33-41 (emphasis added); see also 5:55-57 (same). NetApp also relies on the preamble to claim 9, arguing that the method claim is directed to “a disk memory system” containing “a method of storing data records for at least one associated data processor.” Ho Decl. Ex. A (‘987 Patent) at 23:40, 44-45. NetApp contends that the preamble language dictates that the entire method of storing data in the claimed fashion in claim 9 is not practiced unless the steps are the continuous method of operation for the referenced disk memory system and associated data processor, and the patented method of the ‘987 Patent is not infringed unless the accused system consistently selects and writes to empty logical tracks. Motion at 11. NetApp contends that, cumulatively, this evidence limits the patent to a system that *never* updates parity by *always* writing to empty logical tracks. Because its products do not always write to empty logical tracks, NetApp argues that it does not infringe.

Sun disputes NetApp’s contention that the preamble to claim 9 (which identifies a “disk memory system, having a plurality of disk drives”) serves as a limitation on the claim. Sun argues that claim 9 is a method claim and the preamble includes no steps that could be part of this claim. Sun also argues briefly that it did not rely on the preamble during prosecution to obtain allowance of the claim, the preamble is not mentioned or relied on in the body of the claim and provides no antecedent basis for any term, and the preamble offers no definition of any limitation and instead simply reflects an introduction to the claim. Sun further argues that portions of the patent relied on

---

Therefore, this last piece of evidence has been disregarded, and in any event is unnecessary.

1 by NetApp refer to the design of a “data storage system,” and are separate from the asserted method  
2 claims. Sun contends that nothing in the patent purports to limit the method claims, as opposed to  
3 the system claims. Opp. at 17.

4 With respect to the cited portions of the prosecution history of the ‘987 Patent, Sun argues  
5 that it amended the specification to further explain the disadvantages of the prior art, including that a  
6 parity update was required every time data was modified, and that it “solved” the problem because  
7 parity is not updated when the invention is practiced. Sun argues that its “always” and “never”  
8 statements during prosecution do not address the claims or their scope, and do not support NetApp’s  
9 position that the patent requires cyclical performance but instead explain the benefit that accrues  
10 each time the method steps are performed. Sun maintains that the method claims do not require  
11 repeated performance of the method in the system, but rather that even practicing the method just  
12 once infringes. See Opp. at 15. NetApp persuasively counters that this interpretation is nonsensical  
13 and overreaches because, if it were correct, then any file system that ever writes new or modified  
14 data to an empty logical track, even if only once in the first use of a brand new system, would  
15 infringe. Reply at 1.

16 During oral argument, NetApp cited three cases for the position that the ‘987 Patent  
17 disclaims ever updating parity because a “disclaimer does not require express claim language if  
18 statements in the specification and prosecution history are clear.” See Seachange Int’l Inc. V. C-  
19 Cor. Inc., 413 F.3d 1361, 1372-73 (Fed. Cir. 2005) (where applicant argues that a claim possesses a  
20 feature the prior art does not possess in order to overcome a prior art rejection, argument may serve  
21 to narrow scope of otherwise broad claim language); see also Ormco Corp. V. Align Tech., Inc., 498  
22 F.3d 1307, 1316 (Fed. Cir. 2007) (attributing a broader meaning to claims than that indicated in the  
23 patent and prosecution history would be to “ignore the totality of the facts of the case and exalt  
24 slogans over real meaning.”); Omega Eng. Inc.v. Raytek Corp., 334 F.3d 1314, 1323 (Fed. Cir.  
25 2003) (prosecution disclaimer precludes patentees from recapturing through claim interpretation  
26 specific meanings unambiguously disclaimed during prosecution.”).

27 The Court has previously considered related arguments in connection with claim  
28 construction, and held that:

The express language of these method claims by itself does not compel NetApp's proposed limitation that would require an accused product to always practice these claims in order to infringe. For example, claim 9 is a method claim as recognized by dependent claims 10 through 12 ("The method of claim 9 further comprising . . ."). The claim preamble specifies the configuration of the system in which the method operates, i.e., "[i]n a disk memory system." . . . Even assuming (without deciding) that the preamble is a limitation, see NetApp Suppl. Cl. Const. Br. at 4 n.1, the preamble only describes a system of multiple disk drives configured into two or more redundancy groups each consisting of at least two disk drives, without stating whether the method must be exclusive. Furthermore, the claim 9 preamble does not refer to "[t]he system of claim 1," unlike, e.g., claims 2, 4, or 8. Sun is correct that an accused product need not always practice a method to infringe. See Bell Commun. Research, Inc. v. Vitalink Commun., 55 F.3d 615, 622-23 (Fed. Cir. 1995).

Even though the claim language by itself would not compel NetApp's interpretation, there is strong support for it in the specification and prosecution history. The specification makes clear that the invention itself, unlike the prior art, never updates parity data in the redundancy group and avoids the negative impact on performance in the prior art, using strong and unequivocal language:

"A performance improvement is obtained by *eliminating redundancy data updates* in the redundancy group by writing modified virtual track instances into previously emptied logical tracks and marking the data contained in the previous virtual track instance location as invalid. Logical cylinders containing a mixture of valid and invalid virtual tracks are emptied by writing all the valid virtual tracks into a previously emptied logical cylinder as a background process." '987 patent, Abstract (emphasis added);

"This system avoids the parity update problem of the prior art by *never* updating the parity in a data redundancy group." Id. at 3:33-37 [Solution] (emphasis added);

"Thus, *all* redundancy data creation, writing and free space collection occurs in background, rather than on-demand processes. This arrangement avoids the parity update problem of existing disk array systems and improves the response time versus access rate performance of the data storage subsystem by transferring these overhead tasks to background processes." Id. at 5:21-28 [Solution] (emphasis added);

"The data storage subsystem of *the present invention* uses a plurality of small form factor disk drives in place of a single large form factor disk drive to implement an inexpensive, high performance, high reliability disk drive memory that emulates the format and capability of large form factor disk drives. This system avoids the parity update problem of the prior art by *never* updating the parity. Instead, all new or modified data is written on empty logical tracks and the old data is tagged as obsolete." Id. at 5:55-59 [Detailed Description of the

Drawing] (emphasis added).

NetApp also relies on excerpts from the prosecution history to demonstrate that the applicants made a clear and unambiguous disavowal of claim scope by distinguishing the invention from the prior art disk array systems as “*never* updating parity in any redundancy group at any time.” NetApp Br. re Pros. Hist. at 2-3 (quoting ’987 Prosecution History, Sept. 23, 1991 Amendment at 7 (NAB0013338)) (emphasis added). The applicants’ statements were made in response to the examiner’s objection to the specification and rejection of the claims noting that the written description was “unclear why writing data to a new location is an advantage. This simply changes which parity bits need to be regenerated.” *Id.* at 3 (NAB0013334). While the applicants’ statements in the prosecution history arguably do not by themselves amount to a clear and unambiguous statement that a system must always practice the claimed method in order to practice the patent, they lend further support to the strong language in the specification describing the feature of the invention as a whole, and thus may well limit the scope of the invention to require always writing only to an empty track.

12/22/08 Claim Construction Order at 8-10 (emphasis added).

Having further considered the issue in connection with this motion, the Court now finds that the applicants’ unequivocal statements in the prosecution history, made to distinguish prior art, in conjunction with the strong specification language, amount to a clear and unambiguous confirmation that the ’987 Patent requires that the claimed method always be performed, and thus does not encompass an accused system that only practices the claimed method some, or even most, of the time. The word “never” is one of the strongest and most unequivocal words in the English language, and Sun used it (and its equally strong antonym, “always”) repeatedly in explaining its invention in the specification and prosecution history. Shakespeare famously used the word’s chilling absoluteness to great effect during King Lear’s lament over the death of his daughter Cordelia: “Thou’lt come no more, Never, never, never, never, never.”

Indeed, Sun expressly limited its invention during prosecution to distance itself from prior art in response to the PTO’s rejection, and cannot disavow that limitation now. Sun specifically touted the features of *always* writing full stripes to empty logical tracks, and *never* updating parity, as improvements over prior art. Sun’s witnesses have confirmed that this is how the invention operates in order to solve the parity update problem. Sun may not disavow this central feature in order to capture a device that sometimes, but not always, practices it, particularly when the whole point of the invention as recited to the PTO was to achieve the result every time by never updating parity.

1 Additionally, while the body of claims 9 and 57 do not expressly incorporate the preamble phrase  
2 “disk memory system having a plurality of disk drives,” the language of the preamble does appear to  
3 provide “antecedent basis” for the body of the claim.

4 As discussed below, there is no factual dispute that WAFL generally writes full stripes of  
5 data on empty logical tracks some, but not all, of the time. There is, however, a factual dispute  
6 about the frequency that WAFL does so, whether as high as 99% of the time or less frequently. The  
7 real question before the Court on summary judgment is whether this dispute about how much less  
8 frequently than *always* WAFL does so is sufficient to raise a triable issue of fact as to infringement.

9 **2. It Is Undisputed that WAFL Writes to Empty Logical Tracks Some, But**  
10 **Not All, Of the Time**

11 NetApp contends that, in contrast to the patented invention which was designed to solve the  
12 parity update problem by always writing to empty logical tracks, WAFL was designed with two  
13 competing goals in mind: (1) optimizing the efficiency of writing to a RAID array by writing to as  
14 many blocks as possible within a given RAID stripe; and (2) optimizing the efficiency of subsequent  
15 reads by preferentially writing the blocks of a given file on the same disk. Ho Decl., Ex. H (Strange  
16 Depo.) at 335-336, 347. NetApp argues that WAFL’s “write allocation process is agnostic with  
17 respect to whether the RAID stripes it allocates for writing are empty,” because it does not target  
18 empty stripes, but instead targets sets of empty blocks. See Ho Ex. I (Mattson Expert Report) at ¶  
19 60; see also Ex. J (Haeberli Expert Report) at ¶ 10. NetApp contends that WAFL often allocates  
20 space in which there are no empty stripes. See Ho Dec. Ex. H (Strange Depo) at 349 (discussing  
21 fact that in the typical case in an allocation area, there are some blocks that are already allocated so  
22 it is not possible to write to it sequentially).

23 According to NetApp, WAFL does not always selectively write to empty RAID stripes, and  
24 when it does so “it is merely a side effect of selecting a storage area in a RAID group that is  
25 relatively under-allocated.” See Ex. G (Strange Depo.) at 345-346. It analogizes its selection  
26 process to a two-dimensional grid where the vertical columns are disks horizontally aligned with  
27 each other, and contends that the ‘987 patent focuses solely on writing to empty horizontal stripes,  
28 while WAFL looks for relatively empty areas on the grid, both horizontal RAID stripes and vertical

individual disks. See Reply at 7:15-22. NetApp further argues that even Sun’s own expert, Dr. Smith, acknowledged that there are occasions when NetApp products do not write full stripes. See Ho Decl. Ex. E (Smith Depo.) at 316-317; Ex. K (Smith Expert Report) at ¶ 87 (stating that “WAFL is designed to write full stripes *when it can*”) (emphasis added). NetApp argues that, because both of its own witnesses and Sun admit that WAFL does not *always* perform the claimed method (i.e., write full stripes to empty logical tracks), it cannot be literally infringing.

Sun does not really dispute that WAFL does not *always* write to empty logical tracks, but instead contends that NetApp’s argument that WAFL is “agnostic” with respect to whether it writes to empty logical tracks is false. According to Sun, WAFL is optimized to write full stripes on empty tracks, and writes to empty logical tracks at least 97.5% to 99.6% of the time. See Opp. at 1:28; see also Sun’s Oral Argument Presentation at Sun2-SJ-987-8 (“NetApp tests confirm full stripes are written virtually all the time”) (emphasis added); see also Sun2-SJ-987-17. Sun relies on, among other things, the testimony of NetApp’s expert, Dr. Strange, who stated that “WAFL attempts to write full stripes when it can.” Williamson Decl., Ex. 1 (Strange Depo.) at 137-138; see also id. at 149-51, 161, 163, 165, 189; Ex. 2 (Strange Depo.) at 344-47. Sun also cites WAFL documentation for this position. See, e.g., Williamson Decl. Ex. 3 (NetApp Topology Cache document) (WAFL understands physical layout of underlying storage aggregate in part because it allows WAFL to try to write full stripes). However, Sun fails to mention that the NetApp Topology Cache document goes on to discuss in detail the reasons that “sometimes it’s not desirable to write a full stripe,” and notes that WAFL may not write full stripes because “WAFL also tries to optimize for sequential reads by placing sequential FBNs next to each other on a single disk,” which may cause WAFL to perform partial stripe writes. Id.

Sun also argues that, in addition to trying to write full stripes, WAFL succeeds in doing so almost all of the time. See Ex. 10 (Smith Depo.) at 317-321 (NetApp benchmarking data shows that WAFL writes to empty stripes at least 97% to 98% of the time); Ex. 11-12 (data from November 10, 2009 test showing that NetApp filer wrote full stripes 98% of the time); Exs. 13-17 (data from similar tests showing WAFL writing full stripes approximately 97.5% to 99.6% of the time). At oral argument, Sun presented evidence of new testing data showing 99.9% and 99.75% full stripe writes.



NetApp responds that the testing data is somewhat misleading, and also proves its point that WAFL does not *always* write full stripes to empty logical tracks. NetApp contends that the data cited above on which Sun relies came from tests performed on completely empty file systems where almost all of the logical tracks started out empty, so the system had to choose empty tracks in almost every instance, at least initially. Reply at 8; see also Ho Reply Decl. Ex 2 (Strange Depo.) at 500. NetApp notes, without citation to evidence other than a Reply Declaration by Mr. Strange with no testing data attached, that tests on non-empty filesystems reveal that such systems write to full stripes less frequently, in some cases less than 50% of the time. Strange Decl. ¶ 3. NetApp also notes, without citation, that some of the benchmark tests relied on by Sun were performed on systems containing 1+1 Raid configurations, which, as discussed below, contain only one block per disk and therefore always select the empty block which is read by the test as a “stripe.” However, at oral argument Sun demonstrated that the large majority of this testing was done on configurations other than 1+1 or 1+2, and most contained 12 or 14 data disks. Further, the RAID groups with only one data disk actually wrote less full stripes per second than the groups with multiple data disks. See Sun2-SJ-987-15, 16. Therefore this point is unhelpful to NetApp. However, NetApp persuasively points out that even the tests relied on by Sun show that WAFL does not always write to empty stripes. See Ho Reply Decl., Ex. 2 at 500.<sup>3</sup>

At oral argument, Sun further argued that the claim as construed does not actually require writing “full stripes;” instead, it requires writing to “empty logical tracks.” Sun attacked NetApp for focusing on evidence that WAFL does not always write “full stripes,” when it deems the relevant question to be whether it writes to “empty logical tracks.” However, this is a distinction without a difference because the undisputed evidence shows that WAFL does not always do either. According to Sun, the testing evidence shows that the accused products are optimized to write to empty logical tracks almost all of the time and only a small percentage are partial stripe writes. Sun contends that

---

<sup>3</sup>NetApp also points to the admission of Sun’s expert that “sometimes the NetApp products do update parity in place.” Smith Depo. at 316-17; Smith Expert Report at ¶ 86 (“If Data ONTAP is updating some number of data blocks in a stripe but not all, then Data ONTAP must read the data from the blocks that are not being updated in order to calculate parity for the stripe.”). This reinforces NetApp’s argument because WAFL only updates parity in place when it has not written to an empty logical track.



1 this creates an inference that these partial stripe writes could all have been to empty logical tracks  
 2 and therefore NetApp products always write to empty logical tracks (even if not a full stripe). To  
 3 support the inference, Sun relies on the Strange Reply declaration for what it omits: the declaration  
 4 does not specifically deny that the few partial stripe writes could have been to empty logical tracks,  
 5 and does not expressly state that the tests were on empty file systems or do not reflect typical  
 6 conditions. While the Strange Reply Declaration is fairly general in terms of its examination of the  
 7 testing data, it does not lend support to Sun's hypothesis that 100% of the partial stripe writes were  
 8 to empty logical tracks. An inference has to be reasonable, and nothing suggests that every partial  
 9 stripe write is to an empty logical track.

10 Sun also points out Dr. Mattson's testimony that, over time, many of the blocks in a WAFL  
 11 system may become full and it becomes more difficult for the write allocator to find a relatively  
 12 empty allocation area. Williamson Decl. Ex. 19 (Mattson Report) ¶ 55. Sun argues that the '987  
 13 Patent does not claim a method of utilizing disks of infinite capacity, and therefore does not address  
 14 what happens in the situation where a system runs out of empty logical tracks. Sun's point appears  
 15 to be that WAFL always writes full stripes until there are no more full stripes left to write so WAFL  
 16 always practices the asserted claims of the patent for most of its product life, but Sun presents no  
 17 evidence of this. Sun also suggests that WAFL's partial stripe writes are generally clumped  
 18 together, rather than evenly distributed throughout the writing, meaning that the majority of the time  
 19 there is an uninterrupted string of full stripe writes. However, the testimony of Dr. Strange that Sun  
 20 cites for this proposition is speculative. See Williamson Decl. Ex. 10 at 320-21. Neither of these  
 21 arguments aid Sun, as they essentially admit that WAFL does not always write full stripes to empty  
 22 logical tracks. Even taking the evidence in the light most favorable to Sun and assuming that WAFL  
 23 does so 99% of the time, 99% is not "always."

### 24 **3. A Reasonable Juror Could Find that NetApp's 1+1 and 1+2 RAID** 25 **Configurations Write Full Stripes to Empty Logical Tracks**

26 In addition to its general arguments about WAFL technology being optimized to usually  
 27 write to empty logical tracks, Sun argues that two types of RAID configurations used by NetApp and  
 28 at least one of its customers always infringe claim 9 by always writing full stripes. Specifically, Sun

1 argues that NetApp filers can be configured to include: (1) RAID 4 groups with only one data disk  
 2 and one parity disk (“1 + 1 RAID group”) or (2) RAID DP groups with one data disk and two parity  
 3 disks (“1 + 2 RAID group”). Williamson Decl. Ex. 1 (Strange Depo.) at 97-98; Ex. 6 (Smith  
 4 Report) ¶¶ 869-70; Ex. 10 (Smith Depo.) at 471-72. Sun argues that, in both of these “one data  
 5 disk” configurations, every write to a block is a full stripe write because there is only one block in a  
 6 given stripe. See Ex. 1 (Strange Depo.) at 204; Ex. 6 ¶¶ 869-70 (citing Strange testimony). The  
 7 testimony of Mr. Strange that Sun relies on for this argument does not entirely support its position.  
 8 While Mr. Strange did say that any writes to a 1+1 RAID group constitute a full stripe write because  
 9 there is only one block, he almost immediately clarified that this configuration is a degenerate case  
 10 and he views it more like “disk mirroring” than a RAID configuration. Id. Ex. 1 at 204. He goes on  
 11 to state that, “I’m not sure I would define . . . a configuration that is ultimately a mirror  
 12 configuration as being a stripe at all. It’s really just a single block.” Id. Ex. 1 at 204-205.  
 13 Therefore, Dr. Strange’s testimony on this point is equivocal at best.

14 More persuasively, however, Sun argues that the preamble of claim 9 of the ‘987 Patent  
 15 specifically recites a “plurality of disk drives being configured into at least two redundancy groups,  
 16 each redundancy group consisting of at least two disk drives,” and therefore expressly encompasses  
 17 this type of one data disk and one parity disk setup. Opp. at 9. Sun notes that NetApp’s  
 18 interpretation discounting these configurations as “degenerate” would require three disks, contrary  
 19 to claim 9’s “at least two disk drive” preamble. Sun further argues that the Court’s construction of  
 20 “first available memory space” does not require “stripes” or the use of more than one data disk, and  
 21 that NetApp’s own testimony establishes that it considers 1+1 RAID groups to have “stripes.” See  
 22 Williamson Decl., Ex. 21 at NA0526908; Ex. 1 at 204. Sun also points out that in prior art, if a data  
 23 block was modified, the changes took place in the same block and parity would need to be updated,  
 24 but in a 1+1/1+2 configuration updates are written to a different block and not overwritten. Sun  
 25 argues that the reason NetApp’s customer, Earthlink, uses this configuration is to avoid consistency  
 26 point reads (which include updating parity in place and are part of the “performance bottleneck” the  
 27 invention was designed to avoid). Ex. 21 at NA0526908 (“Striping writes across a number of 2 disk  
 28 raid groups within a volume eliminates tetris reads completely.”). Finally, Sun points out that

1 NetApp's expert, Dr. Mattson, does not opine that these configurations do not satisfy the "first  
2 available memory space" limitation. Paragraph 53 of his report comes the closest, but does not  
3 specifically address the 1+1/1+2 configurations so is not helpful to NetApp.

4 In response, NetApp concedes that there is no factual dispute that these configurations will  
5 always write to an empty single block, but defines the issue as whether writing to a single block  
6 constitutes writing a full stripe. Reply at 8. NetApp correctly argues that the parties have both  
7 previously asserted that a RAID "stripe" refers to the corresponding locations of *multiple* disks of a  
8 redundancy group. See Dkt. No. 60 (NetApp Opening Claim Construction Brief) at n.2; Dkt. No. 69  
9 (Sun's Responsive Claim Construction Brief) at 1-2. NetApp contends that equating writing a single  
10 block to writing a full stripe is akin to arguing that a method of optimizing seat assignments by using  
11 full rows would also be infringed whenever chairs are not arranged in rows but instead lined up one  
12 behind another, because in that case each seat could be considered a row of a single chair. See  
13 Reply at 9. While this analogy has some logical appeal, NetApp does not persuasively address the  
14 fact that its own documentation and testimony shows that NetApp internally has referred to writes  
15 on these 1+1 and 1+2 RAID configurations as "stripes." NetApp also never addresses Sun's points  
16 that the patent itself, and the Court's construction, do not require more than one data disk or the use  
17 of "stripes," nor does NetApp provide any authority to ignore a "degenerate case" that practices the  
18 claims.

19 At oral argument, NetApp further argued that the patent makes clear that there are multiple  
20 blocks on a track, but did not persuasively explain how. NetApp cites a passage in the specification  
21 which states: "Each physical track is combined with N-1 other physical tracks to form the N data  
22 segments of a logical track." Ho Decl. Ex. A ('987 Patent) at 5:10-12. NetApp also cites the  
23 specifications and description that: "Control circuitry reconstructs the data stored on each physical  
24 track of the failed disk drive, using the remaining N-1 physical tracks of data plus the associated M  
25 physical tracks containing redundancy segments of each logical track." Id. at 4:26-30. However,  
26 NetApp did not explain how these passages rule out a track with just one data disk. NetApp relies  
27 primarily on the purpose of the invention, contending that the 1+1/1+2 configurations do not update  
28 parity in place because they are essentially mirrored disks, so there is no risk of the parity update

1 problem and the “solution” contemplated by the ‘987 Patent is inapplicable. While this argument  
 2 has considerable appeal and may ultimately persuade a jury, it is not sufficiently tethered to federal  
 3 circuit authority and the language of the patent language itself to justify granting summary judgment  
 4 as to these configurations. Sun is correct that the patent itself does not require “stripes” and the  
 5 Court’s construction also does not require “stripes.”

6  
 7 **4. Because the Patent Requires that the Claimed Method Always Be**  
 8 **Performed, Caselaw Supports Non-Infringement Except As To the 1+1**  
 9 **and 1+2 Configurations**

10 NetApp analogizes this case to Ferguson Beauregard/Logic Controls, Division of Dover  
Resources, Inc. v. Mega Systems LLC, 350 F.3d 1327 (Fed. Cir. 2003) for the position that, where a  
 11 claim requires that a certain method must always be performed, it will not be infringed if that  
 12 method is only sometimes performed. In Ferguson, the Federal Circuit rejected the patentee’s  
 13 argument, which relied on Bell Communications Research, Inc. v. Vitalink Communications Corp.,  
 14 55 F.3d 615 (Fed. Cir. 1995), that “an accused product that sometimes, but not always, embodies a  
 15 claim nonetheless infringes. Id. at 1346. The Court reasoned that the claim at issue required the  
 16 consistent use of a recited adjustment step, whereas the accused device performed the adjustment  
 17 step only occasionally. The Court in Ferguson distinguished Bell on the basis that, there, the  
 18 accused product sometimes accomplished the entire method of the claim, whereas in the case before  
 19 it the claimed method required that an adjustment step always (not sometimes) be performed. Id.  
 20 NetApp cites other cases in which courts have declined to follow Bell where infringement was not  
 21 possible by only sometimes performing the steps of the claim. For example, in Sunrise Medical  
 22 HHG, Inc. v. AirSep Corp., 95 F. Supp. 2d 348 (W.D. Pa. 2000), the accused device sometimes  
 23 supplied a patient with a dose of oxygen while the patient exhaled. The court rejected an argument  
 24 for infringement based on Bell because, though it was rare, the accused device was capable of  
 25 delivering such a pulse, something the patent specification and prosecution history stated should  
 26 never occur.

27 Sun argues that Bell applies here and establishes that, as long as the series of steps set forth  
 28 in the method claim are practiced even once, infringement has occurred. Opp. at 11. The Bell

1 patent was directed to connecting local area networks of computers or telephones by transmitting  
2 message packets, and the asserted claim in Bell recited a “method for transmitting a packet over a  
3 system comprising a plurality of networks.” Bell, 55 F.3d at 618. Sun points out that, despite the  
4 fact that the patent indicated that innumerable packets would be transmitted on an ongoing basis  
5 over a typical system, the method claim at issue recited a series of steps relating to transmission of  
6 an individual packet. Id. When each of the steps was performed in connection with an individual  
7 packet, the Court found infringement even though the steps might not always occur. Id. Sun argues  
8 that this case is similar, in that claim 9 is directed to “a stream of data records” and “modifications  
9 made to one of said data records” and recites a series of steps that must occur when a stream of data  
10 is received and modified, but does not purport to require the cyclical performance of the method’s  
11 steps in order to practice the method. Opp. at 14. Sun argues that the same is true for the other  
12 asserted claims. See Opp. at n. 6.

13 At oral argument, Sun attempted to distinguish Ferguson on the basis that there the claims  
14 themselves recited repeated cyclical performance, i.e., “for the next cycle of the well” and “each  
15 successive cycle,” whereas the claims at issue do not mention “cycle.” Sun further argued that the  
16 claims themselves do not require exclusive performance, which is true as to the body of the claims,  
17 but not the preamble, which indicates that the method is directed to a system. Sun also argued that  
18 “neither the specification or the prosecution history state that the claimed method must be practiced  
19 to the exclusion of other methods,” because the statements relied upon by NetApp are directed to the  
20 problem of the prior art (requiring parity update every time data modified) and the benefit of the  
21 invention (not updating parity every time) – not the claims themselves. However, and as discussed  
22 above, the claims are directed to a method that results in never updating parity when the invention is  
23 practiced, so the stated benefit would not accrue unless the method was practiced every time.

24 Sun also cites Praxair, Inc. v. ATMI, Inc., 543 F.3d 1306, 1325 (Fed. Cir. 2008) for the  
25 proposition that it is generally inappropriate to limit claims to exclude devices because they do not  
26 serve a purpose of the invention. In Praxair, however, the Federal Circuit found that providing  
27 “uniform” capillaries was only one of multiple stated purposes of the overall invention, and the  
28 specification and claims demonstrated that it was not intended to be a requirement of the

independent claim in question. In contrast, here writing to empty logical tracks in order to never update parity is the stated purpose of the invention, and the solution to the problem is expressly singled out in the specification and file history. Sun also cites E-Pass Technologies, Inc. v. 3COM Corp., 343 F.3d 1364, 1370 (Fed. Cir. 2003), which Praxair relied on, for the position that, “The court’s task is not to limit claim language to exclude particular devices because they do not serve a perceived ‘purpose’ of the invention.” In E-Pass, also in the claim construction context, the Federal Circuit rejected a district court’s reasoning that a bulky credit card could not be within the claim language because it would not serve the purpose of the invention to “simplify the use of credit cards.” The court noted that, “[a]n invention may possess a number of advantages or purposes, and there is no requirement that every claim directed to that invention be limited to encompass all of them.” Id. at 1370. E-Pass differs in that there it appears that the district court was making a subjective judgment call as to what feature would or would not simplify the use of credit cards in the context of multiple possible purposes of the invention, concluding that an irregular shape would not. Here, the claim itself requires writing to empty logical tracks, and it is undisputed that the reason for this requirement is to advance the central purpose of updating parity, whereas the accused device does not always do so but instead sometimes updates parity

Thus, Ferguson and the other cases cited by NetApp are more on point here. Because there is no factual dispute that WAFL generally does not *always* perform the claimed method, summary adjudication of non-infringement on this ground is granted, with the exception of the 1+1 and 1+2 configurations, because a reasonable juror could find that they always write to empty logical tracks as claimed by claim 9.

## 5. Doctrine of Equivalents

Neither party has devoted significant effort to pursuing distinct arguments for or against the doctrine of equivalents in connection with this Motion beyond those already made regarding literal infringement, perhaps in recognition of the fact that “never” and “always” have no equivalent in “occasionally” or “usually.” NetApp points out again that WAFL uses a fundamentally different approach for selecting disk space for writing data and fails to accomplish the stated goal to always write data to empty logical tracks in order to avoid updating parity, and that Sun distinguished the

1 prior art on this basis. NetApp also persuasively argues that the patent and its prosecution history  
2 explicitly teach that there is never a parity update, so Sun cannot reclaim a system that sometimes  
3 updates parity under the doctrine of equivalents.

4 Sun reiterates that WAFL was specifically designed to write to fill stripes when possible and  
5 avoid parity updates, and therefore serves the same goal of avoiding the parity update problem. Sun  
6 does not respond to NetApp's prosecution history estoppel argument in the context of the DOE,  
7 though presumably its position is similar to that described above – that the statements relied on by  
8 NetApp were “merely” made in response to the examiner's concerns about the benefit of the  
9 invention over prior art and therefore can be disregarded.

10 Sun's doctrine of equivalents theory fails for essentially the same reasons as direct  
11 infringement, based on the undisputed fact that WAFL does, at least sometimes, update parity in  
12 place. Further, Sun explicitly disclaimed this method during patent prosecution. Under these  
13 circumstances, the Court finds that “sometimes,” “usually,” or even “virtually always” writing full  
14 stripes to empty logical tracks cannot be the equivalent of “always” writing full stripes to empty  
15 logical tracks. However, as explained above, there is some evidence that the 1+1 and 1+2  
16 configurations always write full “stripes” to empty logical tracks. Therefore, summary adjudication  
17 of this issue is warranted except as to these two configurations.

#### 18 E. The “In Response To” Limitation

19 writing, **in response to the subsequent receipt of modifications to one of said**  
20 **data records** stored in said first available memory space in said selected  
21 redundancy group from said associated data processor, said modified data record,  
22 exclusive of the rest of said received stream of data records and said redundancy  
23 data associated with said received stream of data records written in said first  
available memory space, in second available memory space in one of said  
redundancy groups by including said modified data record with a stream of data  
records subsequently received for said step of writing

24 NetApp argues that summary judgment is also appropriate because it does not literally  
25 infringe the “in response to” limitation of the asserted claims, especially with respect to the  
26 “writing” step. NetApp relies on this Court's construction of the term “in response to” in the context  
27 of U.S. Patent Number 5,549,857, where it construed the term to mean “after and in reaction to.”  
28 See Ho Decl. Ex. F (9/10/08 Claim Construction Order) at 17. NetApp further relies on this Court's



1 discussion of that construction in connection with summary judgment, where the Court noted that it  
2 intended the “in response to” limitation to be a “real, cause in fact trigger.” Id. Ex. G (11/16/09  
3 Order Granting Summary Judgment) at 8.

4 Sun contends that this entire argument is procedurally improper because it violates the  
5 Court’s Order regarding this round of summary judgment motions. See Dkt. No. 121 at 1 (“The  
6 Court will first consider summary judgment motions that involve issues relating to the claim terms  
7 that the Court has already construed.”). The Court did specifically order that summary judgment  
8 motions relate to previously construed claim terms, and the Court has not construed “in response to”  
9 in the context of the ‘987 patent. NetApp counters that the term was construed in a related case, the  
10 two patents concern the same subject matter, the inventors are the same, and similar phrasing is used  
11 in the claim language of both patents, so no further construction of the term is necessary. See Reply  
12 at 11. However, the construction of the same term in a related patent does not necessarily carry  
13 over.

14 Sun initially argued that supplemental briefing should be provided before the Court can  
15 decide this dispute. However, at oral argument, when questioned about the need for additional claim  
16 construction briefing, Sun conceded that “after and in reaction to” was an appropriate construction  
17 of the term for this patent also, but disagreed whether the accused technology infringes when so  
18 construed. Both parties briefed the issue and argued it as if no further construction was necessary.

19 Substantively, NetApp’s Motion distinguishes the patent as requiring that a second “writing”  
20 step be performed with *each* modification of one previously written data record, while, in contrast,  
21 WAFL does not act on each write request as it is received but rather the data is stored and later  
22 written when a consistency point is performed. Motion at 14. NetApp argues that there are many  
23 different conditions that could cause a consistency point (which results in a “write”) to be initiated in  
24 WAFL – such as scheduled performance at a given interval, when there are too many dirty buffers,  
25 or when a snapshot is created – but the “subsequent receipt of modifications to one of said data  
26 records” is never the real, cause-in-fact trigger. NetApp contends that Sun’s expert Dr. Smith did  
27 not cite evidence that a consistency point write is ever triggered by modification of one previously  
28 written data record, but instead testified that it is the cumulative stream of data records since the last

1 consistency point that triggers the new consistency point. See Ho Decl. Ex. E at 286, 287  
 2 (comparing the last data record in the stream that triggers the consistency point to the straw that  
 3 broke the camel’s back, but noting that “obviously the back was broken by the cumulative amount of  
 4 straw”).<sup>4</sup>

5 Sun counters that there is no evidence that WAFL never checks whether any individual write  
 6 triggers a consistency point, and instead Dr. Mattson testified that WAFL “appears to be  
 7 continuously looking to see, among other things, if the NVRAM is nearing half full and when it is, it  
 8 considers starting a CP.” Williamson Decl. Ex. 19 at ¶ 67; see also Ex. 25 (Haeberli Depo.) at 63  
 9 (when a message is received, the routine in question calls a function which may trigger a consistency  
 10 point if certain conditions are met). Sun argues that this testimony confirms Dr. Smith’s analysis  
 11 that “receipt of every client message, such as a write request, always triggers a consistency point  
 12 check.” Opp. at 23 (citing Ex. 6 (Smith Report) at ¶¶ 72-74). Sun’s argument has two problems.  
 13 First, Sun mischaracterizes Dr. Smith’s testimony in that he does not say that receipt of things such  
 14 as write requests “always” trigger consistency point checks. Second, it is irrelevant whether an  
 15 event triggers a consistency point “check;” what matters is whether it triggers an actual consistency  
 16 point resulting in a write. Therefore, this evidence is not material.

17 Sun also argues that Dr. Smith’s testimony and report confirm that some of the received  
 18 write requests that trigger a consistency point are modifications to a previously written data record.  
 19 Sun cites Dr. Smith’s report at Exhibit 6 paragraphs 72-75 and 176-78, but only paragraph 178  
 20 mentions a consistency point in connection with the receipt of a modified data record, and it  
 21 specifically notes that this receipt would only result in a new consistency point if the modified data  
 22 record resulted in too many dirty buffers. There, the dirty buffers, based on an accumulation, rather  
 23 than the one receipt, cause the new consistency point to occur. Dr. Smith’s deposition testimony at  
 24 Exhibit 10 pages 306-307, relied on by Sun, also does not support Sun’s position. Instead, he admits  
 25 that there are a lot of receipts of data records that do not trigger a write because they do not trigger a  
 26 consistency point, and that some writes trigger consistency points while others simply get stored in

---

27  
 28 <sup>4</sup>Sun argues that this testimony is mischaracterized because it pertains to the “selecting” step,  
 not the writing step. However, the testimony relied upon relates to what triggers a consistency point  
 in WAFL and is not specific to any patent claim term. Therefore NetApp’s reliance on it is valid.

1 cache to be part of a future consistency point.

2 NetApp persuasively argues that Sun has not shown that a consistency point is triggered “in  
3 response to” the receipt of a modified data record. The only evidence Sun puts forth on this point  
4 shows that it is not the receipt of the modified data record, but the fact that it resulted in excessive  
5 dirty buffers putting the system over threshold for a consistency point, that triggered the write. See  
6 Williamson Decl. Ex. 10 at 306 (Q: “There is a lot of receipts of data records that don’t trigger a  
7 write, because they don’t trigger a consistency point, correct?” A: “That’s correct.”). Even without  
8 further claim construction or a requirement that the writing be the real, cause in fact trigger for  
9 transmission of data to a secondary system (as was found in the ‘857 Patent), the undisputed facts  
10 show that the receipt of one modified data record is not what causes the consistency point; i.e., it is  
11 not “in response to” (or “after and in reaction to”) the receipt.

12 Sun also argues that the patent itself encompasses a WAFL-type operation, because nothing  
13 in the patent requires a write operation “each” time a modified data record is received or that the  
14 write operation take place immediately following the receipt of the modification. Sun argues that  
15 the claim language “simply requires a write operation at some later time in response to receiving a  
16 modified data record.” Opp. at 21. Sun contends that the specification teaches accumulating  
17 modified data records in the cache and then writing them, just like WAFL. Ho Decl. Ex. A (‘987  
18 Patent) at 5:1-6 (“The host then modifies some, perhaps all, of the records on the virtual track. Then  
19 as determined by cache replacement algorithms such as Least Recently Used, etc, the modified  
20 virtual track is selected to be destaged to a redundancy group.”), 7:9-22 (similar), 14:7-9 (“In this  
21 system, staging operations are similar to staging in other cached disk subsystems but destaging  
22 transfers are collected into groups for bulk transfers.”). Sun argues that nothing in the specification  
23 says that there can not be more modifications before the write takes place.

24 Sun also relies on Figure 7 (which illustrates the process of the data write operation) and its  
25 description in the patent to describe the claimed method. According to Sun, the receipt of a  
26 modified data record is followed by two intervening steps (703 and 704) before the modified data is  
27 even *scheduled* to be written. Ho Dec. Ex. A at 19:44-48; 20:9-21. Sun argues that the scheduling  
28 process then determines when the modified data record gets written to disk. See id. at 7:17-22;

20:23-50. The write does not take place until several other steps are performed. According to Sun, this process is analogous to WAFL's process of storing modified data and periodically flushing it to disk when performing a consistency point. However, WAFL does not even schedule a write upon receipt of a modified data record, so this argument is not persuasive.

NetApp disputes Sun's position that the '987 Patent itself allows for accumulating modifications and later writing them to disk by relying on Dr. Smith's testimony relating to Figure 7. Dr. Smith testified that Box 704 of Figure 7 indicates the receipt of modifications to a data record (the "trigger"). See Ho Decl. Ex. E at 365-66. He further testified that the "write" occurs in the subroutine referenced by Box 705 (labeled in the Figure as "schedule modified virtual track to be written"), and specifically takes place in Box 709. Id. at 366; see also Ex. A. at 20:39-42 ("Once a free logical cylinder is available, either being the presently open logical cylinder or a newly selected logical cylinder, then at step 709 the control unit 101 writes the updated virtual track instance into the logical cylinder"). NetApp argues that Sun's interpretation of the claim specification, placing the "write" earlier in the process and triggering only "scheduling" of the write, is not supported by Sun's expert testimony.

Sun's response that Dr. Smith's testimony supports its own, rather than NetApp's, position is weak. Dr. Smith specifically identifies the "trigger" as Box 704, and the write as taking place in Box 709 during what is labeled as the scheduling subroutine. This testimony, and the patent itself, show that as soon as an empty cylinder is identified to write to, the write occurs. This is far more proximate in time and causation than Sun's attorney argument would have it, and is not just "at some later time," as Sun argues in its brief.

Given the foregoing, there is no triable issue of fact as to this issue. There is no evidence that the receipt of a modified data record is what triggers a consistency point (other than as an incidental result of that data record creating one too many dirty buffers and putting the system over threshold). And while Sun attempts to rely on its belated characterization of a figure in the patent specification to argue that modified data is received long before a write takes place, even its own expert testifies otherwise, including as to that part of the specification. Therefore no reasonable juror could find that WAFL literally satisfies this claim and the motion is granted on this basis.

1           **F.       The “Exclusive Of” Limitation**

2           writing, in response to the subsequent receipt of modifications to one of said data  
3           records stored in said first available memory space in said selected redundancy  
4           group from said associated data processor, **said modified data record, exclusive**  
5           **of the rest of said received stream of data records and said redundancy data**  
6           **associated with said received stream of data records written in said first**  
7           **available memory space**, in second available memory space in one of said  
8           redundancy groups by including said modified data record with a stream of data  
9           records subsequently received for said step of writing.

10          NetApp contends that WAFL does not literally infringe the ‘987 Patent because  
11          there is no evidence that WAFL writes only one modified data record, “exclusive of” the  
12          rest of the stream of data records. When questioned about this claim term, Dr. Smith  
13          testified that it could be reasonably interpreted as “at least one” and “could be more than  
14          one.” Ho Decl. Ex. E (Smith Depo.) at 361-363. NetApp argues that this interpretation  
15          directly conflicts with the relevant claim language.

16          NetApp concedes that a portion of the specification refers to the host processor  
17          modifying “some, perhaps all, of the records on the virtual track” before the modified track  
18          is written to a redundancy group, and that this language would appear to support Sun’s  
19          position that “one” means “at least one.” However, NetApp points to the prosecution  
20          history, where the examiner suggested amending the claims to “more particularly recite  
21          appellant’s invention” by “clearly recit[ing] that the ‘means responsive to the subsequent  
22          receipt of modifications’ only read/modify/write the data record to be modified, and not the  
23          other data records of the data stream with which it was originally written, or the parity data  
24          record which was originally calculated for the received stream.” See Ho Dec. Ex. P at  
25          NAB0013357; see also Ho Reply Decl. Ex. 3 NAB0013375; NAB0013362; NAB0013369;  
26          NAB0013373 (examiner addition of “the rest of” in the phrase “exclusive of the rest of the  
27          said received stream”). NetApp argues that the addition of the “exclusive of” phrase and  
28          the separate addition of the “rest of” phrase by the Examiner in order to allow the claim  
trumps this portion of the specification and confirms that “one” means “only one,” and no  
other data record or modified data record from the same stream can be written at the same  
time. By contrast, a consistency point flushes all of the records at once. NetApp points out

1 that Sun has presented no evidence that WAFL ever performs a write, triggered by the  
2 receipt of modifications to a data record that contains only one data record, and excludes  
3 other previously written data records.

4 Sun does not appear to dispute that the record being written cannot include the rest  
5 of the unmodified stream of data records. The Court agrees with NetApp that the plain  
6 meaning of “exclusive of” is nailed down by the prosecution history in which the claim  
7 was amended as the Examiner proposed. So NetApp is correct as to the construction of  
8 this term.

9 Sun instead contends that there is a triable issue of fact because NetApp’s expert  
10 does not opine that this limitation is not met, while Sun’s expert does opine that the  
11 limitation is met. See Williamson Decl. Ex. 6 at ¶¶ 76-87, 175-180. Further, Sun argues  
12 that a “modified data record” is different from a data record initially written to the “first  
13 available memory space,” so it is not part of the stream that cannot be written with the  
14 modified data record. According to Sun, under the plain meaning of the claim language,  
15 modified data records can be written to the second available memory space together with  
16 other modified data records because they were never part of the same initial stream.

17 Thus, the dispute boils down to whether Dr. Smith’s testimony supports Sun’s  
18 position, and whether NetApp is correct on the legal issue of whether or not the writing  
19 must be in response to only one modified data record from the stream, or Sun is correct that  
20 the writing can be in response to one modified data record and other unmodified data  
21 records from the same stream. Sun’s expert did state that one can mean “more than one,”  
22 and opined that WAFL meets the limitation because “when a modified data record is  
23 received as part of a write request at a NetApp controller, the modified data record is  
24 written to a new, unused location on disk exclusive of the stream of data records (including  
25 the data record being modified) and associated redundancy previously written to disk.”  
26 Williamson Decl. Ex. 6 at ¶ 177. NetApp has not pointed to the testimony of its own  
27 expert going either way on this point. However, Sun has not put forward any evidence that  
28 NetApp’s products ever perform a write in response to modification of just *one* data record.

1 Moreover, Dr. Smith's testimony runs counter to the proper construction of the terms. The  
2 prosecution history shows that the claim was specifically narrowed to clarify that the  
3 modified data record had to be written alone and without any other part of the stream,  
4 including other modified records. Sun's legal argument that these are not part of the same  
5 stream because they have since been modified is unpersuasive.

6 For the foregoing reasons, no reasonable juror could find that WAFL literally  
7 satisfies this element of the claim. Accordingly, the motion is granted on this basis as to  
8 literal infringement as to all configurations.

9 **G. Doctrine of Equivalents With Respect to "In Response to" and**  
10 **"Exclusive Of"**

11 The parties only briefly argue the issue of the doctrine of equivalents with respect  
12 to the "in response to" and "exclusive of" claim terms, and provided no oral argument on  
13 the issue, perhaps because they recognize that the doctrine of equivalents rises and falls  
14 with the direct infringement theory because it depends on the same issues of claim  
15 construction and prosecution history. As discussed above, a party seeking summary  
16 judgment bears the initial burden of informing the court of the basis for its motion, and of  
17 identifying those portions of the pleadings and discovery responses that demonstrate the  
18 absence of a genuine issue of material fact. Celotex Corp. v. Catrett, 477 U.S. 317, 323  
19 (1986). On an issue where the nonmoving party will bear the burden of proof at trial, as  
20 here, the moving party can prevail merely by pointing out an absence of evidence to  
21 support the nonmoving party's case. Id. If the moving party meets this initial burden, the  
22 opposing party must set forth "specific facts showing a genuine issue for trial." See Fed.  
23 R. Civ. P. 56(e)(2); Anderson, 477 U.S. at 250. If the nonmoving party fails to show that  
24 there is a genuine issue for trial, "the moving party is entitled to judgment as a matter of  
25 law." Celotex, 477 U.S. at 323.

26 NetApp's Motion argues that Dr. Smith did not address the "in response to" and  
27 "exclusive of" claim limitations in his expert report, and therefore NetApp has met its  
28 initial burden of pointing to an absence of evidence. Sun responds that NetApp's




arguments rely on a previously undisclosed claim construction position, but prior to this motion neither party had contended that these terms needed construction, indicating that the correct interpretation of these terms is relatively straightforward. See Ho Decl. Ex. O (8/7/08 Joint Claim Construction Statement) at A-6. Sun also points generally to Dr. Smith's report, including paragraphs 76-87, 101-103, and 176-180, to argue that the function-way-result test is satisfied. Only paragraphs 179 and 180 address the doctrine of equivalents as to these elements, and Dr. Smith's conclusory opinion regarding the equivalency of the invention and the accused device was based on the same erroneous interpretation of the claim language discussed above. Not unlike "never," "exclusive of the rest" is strong exclusionary language ruling out alternative arrangements, and notably this language was strengthened as part of a claim amendment, as set forth above. Sun has pointed to no other evidence supporting its doctrine of equivalents argument. Therefore, Sun has failed to carry its burden of raising a triable issue of fact as to the doctrine of equivalents regarding these limitations and NetApp is entitled to summary adjudication on this issue.

#### IV. CONCLUSION

For the foregoing reasons, the Court concludes as a matter of law that the '987 Patent requires a system that always writes to empty logical tracks to satisfy the "first available memory space" limitation, whereas the undisputed evidence shows that the accused device only "sometimes" or perhaps "usually" does so, except as to the 1+1 and 1+2 configurations where there is a triable issue of fact. However, because no reasonable juror could conclude that NetApp's accused device in any configuration infringes the "in response to" or "exclusive of" claim limitations, either literally or under the doctrine of equivalents, summary judgment of non-infringement of the '987 Patent is GRANTED.

**IT IS SO ORDERED.**

Dated: February 24, 2010

  
ELIZABETH D. LAPORTE  
United States Magistrate Judge